

# Targeted social mobilization in a global manhunt

Alex Rutherford <sup>1,\*</sup>, Manuel Cebrian <sup>3,4,\*</sup>, Iyad Rahwan <sup>1,2,\*†</sup>, Sohan Dsouza <sup>1</sup>, James McInerney <sup>5</sup>, Victor Naroditskiy <sup>5</sup>, Matteo Venanzi <sup>5</sup>, Nicholas R. Jennings <sup>5</sup>, J.R. deLara <sup>6</sup>, Eero Wahlstedt <sup>7</sup>, Steven U. Miller <sup>8</sup>

<sup>1</sup>Computing and Information Science, Masdar Institute of Science and Technology, Abu Dhabi 54224, UAE; <sup>2</sup>School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, UK; <sup>3</sup>National Information and Communications Technology Australia, Melbourne, Victoria 3010, Australia; <sup>4</sup>Department of Computer Science and Engineering, University of California at San Diego, La Jolla, CA 92093, USA; <sup>5</sup>School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK; <sup>6</sup>George Washington University, Washington DC 20052, USA; <sup>7</sup>University of Oxford, Oxford OX1 2JD, UK; <sup>8</sup>Champlain College, Burlington, VT 05401, USA

\*A.R., M.C., and I.R. contributed equally to this work.

†To whom correspondence should be addressed; Email: irahwan@acm.org

**Social mobilization, the ability to mobilize large numbers of people via social networks to achieve highly distributed tasks, has received significant attention in recent times. This growing capability, facilitated by modern communication technology, is highly relevant to endeavors which require the search for individuals that possess rare information or skill, such as finding medical doctors during disasters, or searching for missing people. An open question remains, as to whether in time-critical situations, people are able to recruit in a targeted manner, or whether they resort to so-called blind search, recruiting as many acquaintances as possible via broadcast communication. To explore this question, we examine data from our recent success in the U.S. State Department's Tag Challenge, which required locating and photographing 5 target persons in 5 different cities in the United States and Europe in less than 12 hours, based only on a single mug-shot. We find that people are able to consistently route information in a targeted fashion even under increasing time pressure. We derive an analytical model for global mobilization and use it to quantify the extent to which people were targeting others during recruitment. Our model estimates that approximately 1 in 3 messages were of targeted fashion during the most time-sensitive period of the challenge. This is a novel observation at such short temporal scales, and calls for opportunities for devising viral incentive schemes that provide distance- or time-sensitive rewards to approach the target geography more rapidly, with applications in multiple areas from emergency preparedness, to political mobilization.**

The Internet and online social media are now credited with the unprecedented ability to coordinate the mobilization of large masses of people to achieve remarkable feats that require coverage of large geographical and informational landscapes in a very limited time. Social media has been used to mobilize volunteers to map natural disasters in real-time [1], to conduct large-scale search-and-rescue missions [2], and to locate physical objects within extremely short time frames [3].

Despite the numerous successes attributed to the Internet, mobile communication and social media, we still lack a comprehensive understanding of the dynamics of technology-mediated social mobilization. Open questions remain about essential aspects that determine the success of social mobilization. One such aspect is the relationship between social interaction and geography. Social interaction is an essential driver of recruitment and coordination. However, social interaction is constrained by geography [4], and such constraints exhibit fundamentally different characteristics for large communities [5]. Further, geography is influenced by the nature of the task at hand, as we discuss below.

Consider the task of mobilizing protesters as part of the Occupy Wall Street movement [6]. It has recently been shown that social interaction exhibits a disproportionately high degree of geographical locality, reflecting the movement's efforts to mobilize resources in their local neighborhoods and cities [7].

On the other hand, mobilization for large search-and-rescue operations demands the opposite approach, namely spreading the message and recruiting participants in geographically distant locations. In the *DARPA Network Challenge* (a.k.a. *Red Balloon Challenge*), organized by the Defense Advanced Research Projects Agency, teams

competed to locate and submit the coordinates of 10 tethered weather balloons dispersed at random locations all over the continental United States. The winning team, based at MIT, won the challenge by locating all balloons in less than 9 hours [8]. The team used an incentive scheme to kick start an information and recruitment cascade that resulted in 4,400 sign-ups to the team's Web site within 48 hours. Our earlier analysis revealed that the recursive incentive scheme may have played an important role in maximizing the speed and branching of the diffusion to limits above what is normally observed in viral propagation schemes [9]. Further, data reveals that people managed to recruit acquaintances who are more distant than expected, thus contributing to the rapid coverage of a large geographical area [3].

Another class of mobilization tasks requires geographical propagation that simultaneously spans large distances, while exhibiting targeted spatial dynamics. An example of this is search for a missing person or an object with a known approximate location. Milgram's landmark "small world" experiment showed that people are, in principle, able to find a target individual using 6 hops on the global social network [10]. This result has been reaffirmed in the Internet age in an email-based version of Milgram's experiment [11]. This phenomenon relies on people's ability to form reliable estimates of distance to the target, in order to exploit the large jumps afforded by small world networks as they forward the message to their acquaintances [12–14]. In particular, people rely on *heuristic* information (simple rules of thumb for guiding choice) in the routing of information by the recruitment of acquaintances. Geographical distance, along with non-geographical distance measures – such as similarity of occupation to the target individual – form particularly effective heuristics [15]. For example, if the target is known to be a Professor residing in Kyoto, Japan, one might send it to a friend who lives in Tokyo, Japan, as they are more likely to know someone who lives in Kyoto, who in turn may know someone in academia, and so on.

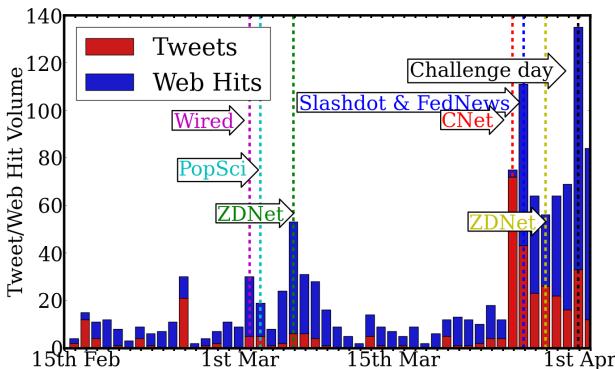
An open question remains as to whether in *time-critical* situations, such as public response to natural disasters, an abduction, or search for a missing child, people are still able to spread information in such a heuristic manner. Humans have a limited amount of time per day to dedicate to social interaction [16], which poses a limit on the effort one can invest in persuading an acquaintance to act. Further, time pressure can affect the way in which people process environmental information [17]. Consequently, people may be expected to resort to so-called *blind* search, focusing simply on the recruitment of as many acquaintances as possible via broadcast messaging [18]. However, while this strategy may be effective at delivering the message to a broad audience, it results in lower effort in finding and mobilizing those recruits that have high *affinity* with the task (due to their location or other characteristics), and are therefore more likely to propagate the message or participate in the required action [19].

We examined the spatial dynamics of global recruitment in the State Department's Tag Challenge, which required competing teams to locate and photograph 5 target "thieves" (actors) in 5 different cities in the US and Europe, based only on a mug shot released at 8:00am local time in each respective city [20]. The targets were only visible for 12 hours, and followed pre-arranged itineraries around the cities of Stockholm, London, Bratislava, New York City and Wash-

ington D.C. Our team successfully located 3 of the 5 suspects [21], winning the competition by remotely mobilizing volunteers through social media using a recursive incentive mechanism that encourages recruitment [22, 23]. This was achieved despite the fact that none of our team members were based in any of the target cities [24].

The challenge provided a rare opportunity to quantify the dynamics of large-scale, global social mobilization in a time-critical scenario from a spatial and temporal perspective. The 12 hour deadline provides a clear urgency. Furthermore, the announcement of the challenge, 2 months in advance, provides a chance to quantify the growth of awareness over time, as we approach the actual day of the challenge, March 31st, 2012. Finally, due to its geographical dispersal over multiple countries and languages, no single small team of acquaintances can conceivably achieve the task without the help of others not directly connected to them. Consequently, people were required to forward messages to acquaintances who are either in the target cities, or whom they believed would be more likely to forward messages towards those cities. Despite the DARPA Network Challenge is very close in aim, it did not provide this opportunity, as there was no information whatsoever about the location of the balloons.

We collected data about the awareness of the challenge, measured by number of hits to the main challenge organizers' Web site, as well as on major social media sites (Twitter and Facebook). We also captured data about the winning team's presence on major social media sites (Twitter and Facebook). This gave us a quantitative view of the growth dynamics of mobilization over time as the deadline approaches. More importantly, by mapping the approximate geographical locations of different social media messages, we were able to quantify the geographical convergence towards the target cities.



**Fig. 1.** Daily volumes of Tag Challenge related Tweets and Web hits up to the challenge day. Major media coverage events are highlighted.

Twitter, the popular micro-blogging service, is an ideal barometer for investigating blind versus heuristic (targeted) mobilization strategies as both modes of communication are available. Users may tweet messages to all their friends (the content is also publicly available if the user chooses this option). Alternatively, a user may mention one or more other users specifically, regardless of whether they are friends or not, by adding the symbol @ followed by the target user's Twitter name. For example, to target a person with user name alex, one simply includes the string @alex in the message. If a tweet is of this second variety, the mentioned user receives a specific alert and is generally obliged to respond, or at least pay more attention to the message. Often, such targeted messaging also leads to subsequent public or private conversations. In the case of the Tag Challenge, such conversations can be seen as an effort exerted on behalf of the recruiter to persuade the recruit to join the cause.

By classifying each challenge-related message to either the broadcast and targeted variety, we were able to investigate the ex-

tent of conscious effort towards targeted mobilization over time as the deadline approaches. In addition, by combining this information with the approximate geographical location of the target audience, it was also possible to investigate whether this targeting was effective in converging towards the target cities geographically.

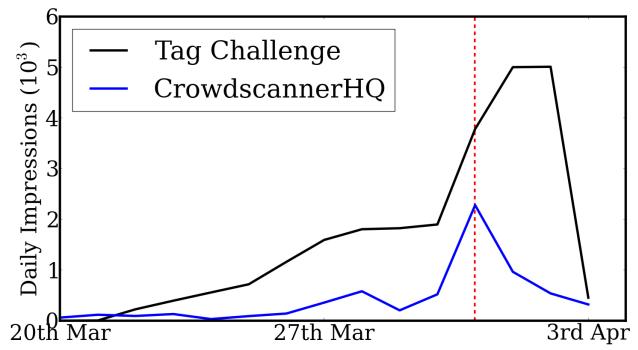
It is important to disentangle two potential explanations of the phenomenon of targeted recruitment in this time-critical social mobilization. One explanation is the explicit effort on behalf of participants to identify and recruit acquaintances who are closer to the target geography. But another explanation is also possible, namely the intrinsic structure of global communication and its role in routing information automatically towards hubs. This is particularly relevant, since two of the target cities, London and New York City, are recognized global hubs. To disentangle the roles played by global communication structure and by individual participant choices, we developed a biased routing model that parameterizes the degree of explicit heuristic targeting, and use it to quantify the behavior observed.

## Results

**Media Exposure.** Fig. 1 shows the daily volume of Tweets related to the Tag Challenge and traffic to the official website (see Materials & Methods). The dates of major media articles concerning the challenge are also indicated. There is clearly some degree of correlation between media coverage and social media traffic. However significant traffic persists on days with no media coverage suggesting that there is also a slower process of peer-to-peer sharing of information about the challenge.

We also see from Fig. 2 that our team's social media presence, measured by the daily number of impressions of our presence on Facebook, provided access to daily volumes of several thousand potential searchers. Although this measure counts repeated exposure by the same users, the total sums to over 29,000. The official Tag Challenge Facebook page also created over 86,000 impressions. We can therefore infer the presence of a hidden network of 'passive recruits' – people who are aware of the challenge, yet are not sufficiently motivated to sign up and recruit others, but who will report sightings of the target. Such a mechanism was found to be a necessary condition for successful social mobilisation in geographical search [25].

**Evidence of Targeted Mobilization.** Fig. 3 shows the distance scaling behaviour of traffic to the Tag Challenge Web site in the 50 days leading up to the challenge. The distance from the originating Internet Protocol (IP) address to the nearest Tag Challenge city was calculated for each unique visitor. After filtering distance independent traffic and smoothing (see Materials & Methods), we observe a strong trend of geographical convergence towards the target cities over time, quantified by the Pearson coefficient ( $r, p$ ) =  $(-0.61, < 10^{-5})$ .

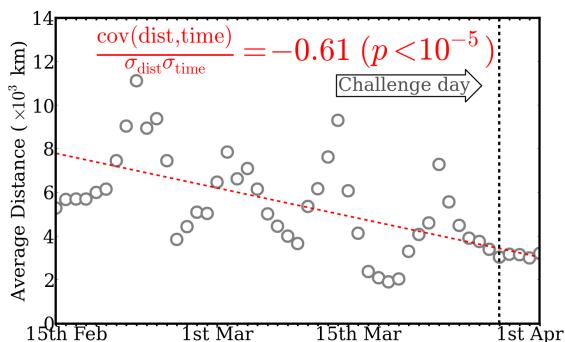


**Fig. 2.** Daily number of impressions on Facebook for the winning team CrowdscannerHQ, and the official Tag Challenge organizers. The vertical dotted line denotes the release of the first mug shots.

Fig. 4 considers the *rate* at which individual users are specifically targeted (i.e. @-mentioned) in the Tweets related to the Tag Challenge. This distinguishes messages which *broadcast* to all followers from those which target *specific* users perceived to be useful for locating the targets (we exclude Tweets from the participating teams from this analysis). The proportion of Twitter traffic targeting individuals increases in the 6 days leading up to the Tag Challenge ( $r, p = (0.825, 0.012)$ ).

This trend is additionally supported by Fig. 5, which considers the *location* of users specifically targeted (@-mentioned) in Tweets. The effect of spurious noise was mitigated with the use of a 4 day moving average. The daily proportion of these targeted users located in the tag cities<sup>1</sup> (with respect to the *total* number of daily targeted users) was seen to increase approaching the challenge day. A strong correlation with time was found ( $r, p = (0.912, 0.002)$  ( $(r, p) = (0.822, 0.012)$ ) using the raw, unsmoothed data). This result suggests that Twitter users successfully route information geographically towards users more likely to locate a target.

The increase in both the *rate* of targeted messaging and its *geographical convergence* suggests that, as time becomes more critical, people become surprisingly *more* rather than *less* targeted in their social mobilization heuristic. This is a novel observation at such short temporal scales (days to hours), and calls for devising viral incentive schemes that provide distance- or time-sensitive rewards to approach the target geography more rapidly, with applications in multiple areas from emergency preparedness [1, 18] to political mobilization [26, 27].



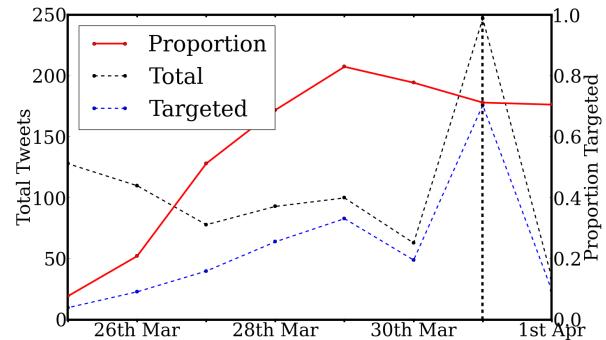
**Fig. 3.** Distance convergence toward Tag Challenge cities of web hits on <http://www.tag-challenge.com>. We consider a moving average of distance filtered daily tweet traffic ( $\text{MA}(\text{prop}^{0.25}(t))_4$ ) (grey circles), which is fit with a linear regression (red line) giving a correlation of ( $r, p = (-0.61, < 10^{-6})$ ).

**Disentangling Targeting Behavior.** The results above suggest the existence of a significant effort by people to mobilize others in a targeted manner, moving towards the target cities. However, it is reasonable to suspect that this observed behavior may be, at least in part, an artefact of the importance of major cities like New York and London — which may receive a disproportionately amount of traffic regardless of the propagation process. Thus it is important to quantify the extent to which we can expect to reach those cities without any deliberate targeting, then use this baseline to quantify the amount of targeting needed to produce the observed behavior in the Tag Challenge.

To investigate this issue, we construct a network of communications between global Metropolitan Statistical Areas (MSA). We use flight frequency data between MSAs as a proxy for social media communication intensity, which have been shown to correlate well (and more strongly than distance) with traffic from Twitter data [28]. Air traffic connections reflect the cultural/linguistic and even post-colonial and post-Commonwealth expatriate ties that have been

found to be present in social networks [29, 30] as well as inter-city economic relations [31] and internet connectivity [32]. An additional advantage of using the air flight network is that we are able to capture the structure of what is a combination of different social media platforms which make up a fragmented global social media ecosystem. This includes not only email but also Facebook, Orkut and Weibo which dominate in North America and Europe, the Lusosphere and China respectively along with many others.

We simulate a random walk over the MSA network, which represents the diffusion of social mobilization using social media and other means of communication (see Materials and Methods for more details). To capture the effect of different mixing of *targeted* and *broadcasting* behaviour, we assign some degree of geographical greediness (targeting)  $g \in [0, 1]$  in making the mobilization decisions. With probability  $(1 - g)$  a random walker on a node chooses to move (i.e. send a message) to a connected node randomly according to the outgoing edge weights (including self-edges capturing local communication within the MSA). With probability  $g$  the walker instead moves greedily to one of its neighbours which enjoys the network-constrained, closest geographic position to any Tag Challenge city (it does this independently of the edge weight). Note that this will generally lead to an *overestimation* of the centralities of the Tag Challenge cities since it assumes that people can successfully leverage any link to a Tag Challenge city no matter how weak it might be. Therefore the degree of greediness (targeting) we report to reproduce our observations should be considered a lower bound. The greedy behavior represents an agent who actively chooses to leverage social ties which are perceived to be more likely to find a target due to privileged location in space [10]. When a walker chooses to move greedily and has more than one Tag Challenge city among its neighbours, it chooses one at random.



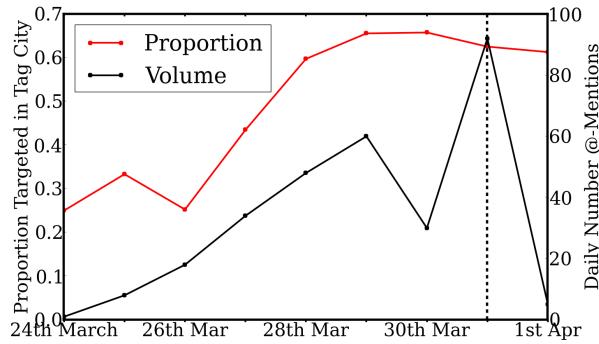
**Fig. 4.** The total daily number of Tweets (black line), the number targeting individuals via @-mentions (blue line) and their proportion (red line). Correlation of targeted proportion with time was found as ( $r, p = (0.825, 0.012)$ )

We perform simulations to determine the stationary probability distributions of the above random walk ( $10^6$  steps per simulation), given various degrees of greedy targeting towards Tag Challenge cities. From this stationary probability we infer the effective centralities of the different cities.

Fig. 6 (red) shows the unbiased centralities without any greedy targeted mobilization. The figure highlights the existence of clear peaks at hubs, including some tag cities themselves. This random walk, corresponding to untargeted broadcast mobilization by participants, leads to 6% of traffic ending up in one of the Tag Challenge cities. While this is a significant proportion in a global network of metropolitan areas, largely driven by the centralities of London and

<sup>1</sup>defined as 25km from the city centre

New York, it is significantly lower than the observed proportion. In particular, as shown in Fig. 4 the proportion of targeted tweets with @-mentions increases to  $\approx 0.7$  as the deadline approaches. The proportion of those tweets that are in one of the target cities is  $\approx 0.65$  (Fig. 5). This means that the proportion of messages reaching the target cities is approximately  $0.7 \times 0.65 \approx 0.46$ , almost an order of magnitude higher than what would be expected by an unbiased, non-targeting random flow of messages.



**Fig. 5.** Daily proportion of @-mentioned users which are located within a tag city. Noise is eliminated by smoothing with a 4 day moving average. Correlation with time reveals a trend given by  $(r, p) = (0.912, 0.002)$

Fig. 6 (bottom left, black) highlights that a significant degree of targeting behavior, corresponding to  $g = 30\%$ , is required to approach the approximate proportion of time spent in the Tag Challenge cities as observed in the data. In other words, people not only need to target others with personalized recruitment messages, but they also need to do so using a geographically informed heuristic at least 30% of the time. Even when restricting the communication network to North America and Europe, to mitigate the affects of linguistic barriers, significant targeting remains necessary to reproduce the observed proportions of traffic. However the diverse originating locations of global traffic to our team’s site suggests that awareness of the challenge did transcend linguistic barriers, justifying consideration of the full global network (see SI Appendix).

## Discussion

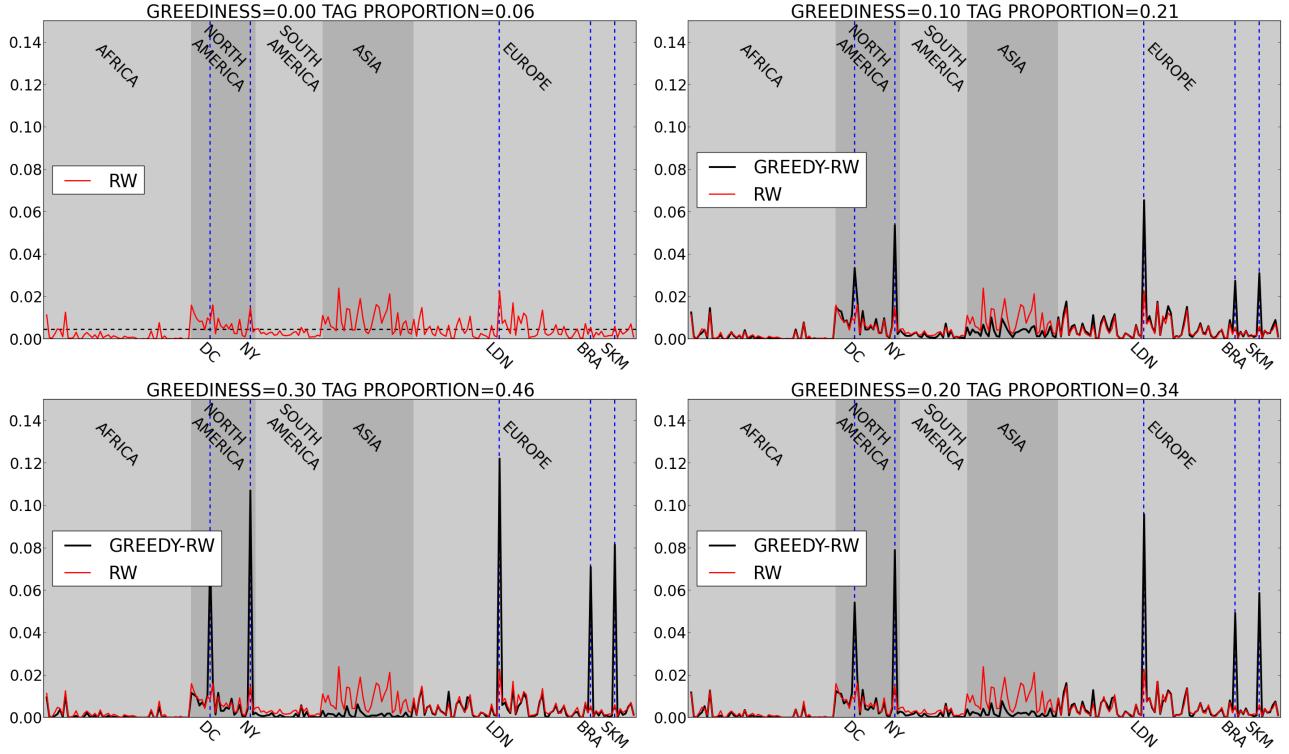
Sixty years ago, social psychologist Stanley Milgram redefined our notion of social distance with his landmark *Six Degrees of Separation* experiment [10], showing that we are, on average, only 6 hops of friendship away from anyone else on earth. Facebook found the

degree of separation to be only 4 in their digital network [33]. Endeavors like the Tag Challenge are set to redefine our conception of the temporal and spatial limits of technology-mediated social mobilization in the Internet age, showing that we can find any person (who is not particularly hiding) in less than 12 hours.

We have shown that this *12 hours of separation* phenomenon relies crucially on the ability of social networks to mobilize in a targeted manner, using geographical information in recruiting participants. The data provides significant support for the presence of geographical targeting, even under time pressure. In fact, we observe that targeting increases as a function of time pressure, as the challenge approaches its deadline.

We were also able to quantify the intensity of targeted mobilization behavior, in comparison with the baseline of untargeted flow of global social media communication. This supports the general notion that social networks are able to tune their geographical communication to suit the task at hand. For example, using Twitter data, it was shown that the Occupy Wall Street social movement in the United States exhibits significant localization (at the state level) when it comes to messages that facilitate resource mobilization and coordination, with reference protest action and specific places and times. In contrast, information flows across state boundaries are more likely to contain framing language to develop narrative frames that reinforce collective purpose at the national level [7]. Our findings complement these results, by contributing towards a general theory that link the purpose of social mobilization to the temporal and spatial dynamics of different forms of communication.

Within high volume social media communications, considerable effort is required to persuade people about the importance of a particular message or cause or even to notice it at all. Both considerations are crucial for a successful mobilisation process. Previous work has shown that shared news stories of interest become obsolete on a timescale  $\approx 1h$  [34] and that the amount of cognitive resources an individual dedicates to online communications is limited and inelastic [35], meaning that the intrinsic importance of the message cannot be relied upon to overcome informational overload and to motivate its sharing. In addition, active interaction with a task requires much more attentional cost to an individual than simple observation [36] and connected individuals vital for propagation also have an associated high inertia [37]. The importance of targeted personal interactions (typified by Twitter @ mentions) can be seen in this context; personalised messages obligate greater cognitive effort from the receiver overcoming the inevitable slide into obsolescence of a single subject over time. Geographical targeting now has an additional advantage beyond the increased chance of recruiting a first hand searcher as the targeting converges; increased personal affiliation of the receiver with the message. The empirical evidence presented above suggests that large distributed communities intuitively understand these considerations and can leverage them in a timely and powerful manner.



**Fig. 6.** Plot of stationary distribution during a random walk on global MSA network, with increasing degree of greediness (targeting) moving clockwise from top left. The red line represents a pure, untargeted random walk, corresponding to pure random mobilization via broadcast messaging. (Top left) The horizontal dashed line represents the uniform distribution of centralities expected in a fully connected graph. The black line in other plots represents a greedy random walk. (Bottom right) When the greediness is increased to 30% we match the observed proportion of targeted messages reaching the Tag Challenge cities. The shading represents MSAs from different continents. The 5 tag cities are marked with vertical, dashed blue lines.

## Materials and Methods

**Twitter.** The Web site Twitter is an extremely popular micro-blogging service which also incorporates a social network. Users create short messages ('tweets') of 140 characters or less which contain text and/or shortened hyperlinks to other webpages or images of interest. Users tweets appear in the *feed* of all other users who have chosen to *follow* her. A user may also opt to make the content of their tweets visible to the public. Tweets contain hashtags to signify that the tweet is relevant to a particular topic i.e. #playTag was a popular hashtag for the Tag Challenge. Users may also choose to target a Tweet to a particular user, regardless of whether the users are connected by a follower/following link, rather than simply broadcasting to her followers. This is done by including a user's Twitter handle e.g. @crowdscannerhq.

We collected the full set of relevant tweets from the period 13<sup>th</sup> February to 10<sup>th</sup> April using a paid service [38] according to appropriate hash tags and keywords or targeted mentions (@ mentions) of competing teams. Tweets originating from @TagTeam\_, @CrowdscannerHQ, @TagChallenge, @Tagteam and @Tag\_Challenge were discarded. Tweets from the participating teams were excluded from these daily totals since the teams had an interest in increasing the daily tweet volumes. The tweets were then manually filtered for relevance by relevant hashtags such as #playTag, #tagchallenge, #tag and any links con-

tained within the tweet. 1263 tweets out of 2181 remained after the filtering process.

Tweets from users with no reliable location information which could be geocoded were discarded, further care was taken to recognise and eliminate artefacts of the geocoding process which led to spurious latitude/longitude coordinates. e.g. 'The world' becoming '(0.0,0.0)'. Tweets originating from within 25km of the defined city centres [39] were considered to originate from the city.

**Facebook.** As the large Web-based social network in the world, Facebook has over 1 billion active users. The daily number of impressions were sourced using the Facebook Insights Application Programming Interface (API) [40]. This covers any user engagement with Tag Challenge page, such as posts on one's "wall" or expressions of approval by friends using the "like" button, etc.

**Google Analytics.** The traffic to the official website was recorded between 14<sup>th</sup> February and 4<sup>th</sup> April. A total of 1000 unique users and their IP addresses were recorded in this period. We used an online service [41] to derive approximate location coordinates from this IP. To mitigate the effect of noise due to the variable volumes of traffic, a moving average was taken for each day, using a sliding window defined as  $(MA(prop^\beta(t))_n = (prop^\beta(t-n) + \dots + prop^\beta(t-1) + prop^\beta(t))/n)$ , where  $prop^\beta(i)$  is the proportion of distance ordered tweets within the  $\beta$ -th percentile on day  $i$  which were within a tag city and  $n$  is the order of the moving average. Fig. 2 corresponds to  $n = 4$  and  $\beta = 0.25$ .

Even the full set of unsmoothed data ( $n = 0, \beta = 1$ ) reveals a geographically convergent trend ( $r, p = (-0.34, < 10^{-5})$ ). We excluded tweets from the Tag teams since the teams may have actively pursued a strategy of geographical convergence skewing the results.

**Simulation.** A coarsened network of air travel connections was constructed as follows. Firstly the largest 220 Metropolitan Statistical Areas (MSA) were considered across all continents. A full list of global airports and connections between them was taken from Open Flights [42]. In order to coarsen the data, airports were agglomerated to the geographically closest MSA using open data [43] [28]. Now the many airports of Greater London; Heathrow, Stanstead, Luton, Gatwick etc are all considered together. This coarsening helps mitigate the effect of anomalous behavior within sparsely populated regional clusters with unusual locality, such as Alaska [44]. We consider the polycentric MSA of Vienna/Bratislava as one single node in the network.

The network edge weights are based on a normalised number of flights between every 2 cities, with self loop weights set to 0.39 representing the probability

of communication within the same MSA [28]. We construct an adjacency matrix representation of the network, namely an  $n \times n$  square matrix  $A$ , where  $n$  is the number of MSAs, and  $A_{ij}$  is the weight of the directed edge between cities  $i$  and  $j$ . The adjacency matrix was row normalised, such that row  $A_i$  represents a probability distribution over the target node reached by a random walker leaving node  $i$ . This results in an adjacency matrix which is nearly symmetric.

We then simulated a random walk over this network. With probability  $g \in [0, 1]$ , so called greediness bias, we move towards the closest Tag Challenge cities. And with probability  $1 - g$  we take a pure random walk with probabilities proportional to the outgoing edge weights. A random walk, with  $g = 0$  corresponds to the eigenvector centrality vector of the different MSAs (see SI Appendix for further details).

**ACKNOWLEDGMENTS.** We would like to thanks David Alan Grier for discussion. We are grateful to all our volunteers that joined the CrowdScanner team.

1. Gao H, Barbier G, Goolsby R (2011) Harnessing the crowdsourcing power of social media for disaster relief. *IEEE Intelligent Systems* 26:10–14.
2. Hellerstein J, Tennenhouse D (2011) Searching for Jim Gray: a technical overview. *Communications of the ACM* 54:77–87.
3. Pickard G, et al. (2011) Time-critical social mobilization. *Science* 334:509–512.
4. Liben-Nowell D, Novak J, Kumar R, Raghavan P, Tomkins A (2005) Geographic routing in social networks. *Proceedings of the National Academy of Sciences* 102:11623–11628.
5. Onnela JP, Arbesman S, González MC, Barabási AL, Christakis NA (2011) Geographic constraints on social network groups. *PLoS one* 6:e16939.
6. Chomsky N (2012) *Occupy* (Penguin Books).
7. Conover MD, et al. (2013) The geospatial characteristics of a social movement communication network. *PLOS ONE* 8:e55957.
8. Tang J, et al. (2011) Reflecting on the darpa red balloon challenge. *Communications of the ACM* 54:78–85.
9. Iribarren JL, Moro E (2009) Impact of Human Activity Patterns on the Dynamics of Information Diffusion. *Physical Review Letters* 103:038702+.
10. Milgram S (1967) The small world problem. *Psychology Today* 61:60–67.
11. Dodds PS, Muhamad R, Watts DJ (2003) An experimental study of search in global social networks. *Science* 301:827–829.
12. Kleinberg J (2000) Navigation in a small world. *Nature* 406:845.
13. Kleinberg J (2000) The small-world phenomenon: an algorithm perspective. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing (ACM)*, pp. 163–170.
14. Adamic LA, Adar E (2005) How to search a social network. *Social Networks* 27:187–203.
15. Watts D, Dodds P, Newman M (2002) Identity and search in social networks. *Science* 296:1302–1305.
16. Miritello G, et al. (2013) Time as a limited resource: Communication strategy in mobile phone networks. *Social Networks* .
17. Ozel F (2001) Time pressure and stress as a factor during emergency egress. *Safety Science* 38:95–107.
18. Watts D, Cebrian M, Elliot M (2013) Dynamics of social media. In Council NR, ed., *Public Response to Alerts and Warnings Using Social Media: Report of a Workshop on Current Knowledge and Research Gaps* (The National Academies Press, Washington, D.C.).
19. Iribarren JL, Moro E (2011) Affinity paths and information diffusion in social networks. *CoRR* abs/1105.3316.
20. (2012), Tag Challenge Press Release. Available online at <http://www.tag-challenge.com/2012/04/04/tag-winner-press-release/>, accessed March 13, 2013.
21. Firth N (2012) Social media web snare ‘criminals’. *New Scientist* .
22. Kleinberg J, Raghavan P (2005) Query incentive networks. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pp. 132–141.
23. Cebrian M, Covello L, Vattani A, Voulgaris P (2012) Finding red balloons with split contracts: Robustness to individual’s selfishness. In *Proceedings of the ACM Symposium on Theory of Computing*, pp. 775–788.
24. Rahwan I, et al. (2013) Global manhunt pushes the limits of social mobilization. *IEEE Computer* 46 (4):68–75.
25. Rutherford A, et al. (2013) The Limits of Social Mobilisation. *Proceedings of National Academy of Science* 110 (16):6281–6286.
26. González-Bailón S, Borge-Holthoefer J, Rivero A, Moreno Y (2011) The dynamics of protest recruitment through an online network. *Scientific Reports* 1.
27. Bond R, et al. (2012) A 61-million-person experiment in social influence and political mobilization. *Nature* 489:295–298.
28. Takhteyev Y, Gruzd A, Wellman B (2012) Geography of twitter networks. *Social Networks* 34:73 – 81, capturing Context: Integrating Spatial and Social Network Analyses.
29. State B, Park P, Weber I, Mejova Y, Macy M (2013) The mesh of civilisations and international email flows. *arXiv* 1303.0045.
30. Ugander J, Karrer B, Backstrom L, Marlow C (2011) The anatomy of the facebook social graph. *CoRR* abs/1111.4503.
31. Beaverstock JV, Smith RG, Taylor PJ (1999) A roster of world cities. *Cities* 16:445–458.
32. Choi JH, Barnett GA, Chon BS (2006) Comparing world city networks: a network analysis of internet backbone and air transport intercity linkages. *Global Networks* 6:81–99.
33. Backstrom L, Boldi P, Ugander J, Vigna S (2011) Four degrees of separation. *arXiv* 111.4570.
34. Wu F, Huberman BA (2007) Novelty and collective attention. *Proceedings of the National Academy of Sciences* 104:17599–17601.
35. Miritello G, Lara R, Cebrán M, Moro E (2013) Limited communication capacity unveils strategies for human interaction. *arXiv preprint arXiv:13041979* .
36. Backstrom L, Bakshy E, Kleinberg J, Lento TM, Rosenn I (2011) Center of attention: How facebook users allocate attention across friends. In *In International Conference on Weblogs and Social Media (ICWSM)*.
37. Hodas NO, Lerman K (2012) How visibility and divided attention constrain social contagion. In *Proceedings of the 2012 ASE/IEEE International Conference on Social Computing and 2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust (IEEE Computer Society, Washington, DC, USA), SOCIALCOM-PASSAT ’12*, pp. 249–257.
38. Hootsuite, Hootsuite. <http://www.hootsuite.com/>, [Online, accessed 9th April 2013].
39. Geohack. <http://bit.ly/Ygy71M>, [Online, accessed 9th April 2013].
40. Facebook insights. <https://developers.facebook.com/docs/insights/>, [Online, accessed 9th April 2013].
41. Ipinfo. <http://ipinfodb.com>, [Online, accessed 9th April 2013].
42. Openflights. <http://openflights.org/data.html>, [Online, accessed 9th April 2013].
43. Geonames. <http://geonames.org>, [Online, accessed 9th April 2013].
44. Guimera R, Mossa S, Turtschi A, Amaral LAN (2005) The worldwide air transportation network: Anomalous centrality, community structure, and cities’ global roles. *Proceedings of the National Academy of Sciences* 102:7794–7799.

Supporting online material for:

Targeted Social Mobilization in a Global Manhunt

## Contents

<b>A Flight Network</b>	<b>1</b>
<b>B Reduced Network</b>	<b>4</b>
<b>C Website Traffic</b>	<b>5</b>

## A Flight Network

We can visualise the adjacency matrix of the MSA network both in terms of raw number of flights (Fig (1)) and the normalised, locality and greediness-adjusted edges (Fig (2)). There are 2 striking features in Figure (1), firstly we see a strong community structure with respect to continents as also observed in [?], particularly within Asia and secondly the high occupation of the diagonal. While the intercontinent connectivity is intuitively understandable, the latter is explained by geography. In regions such as Polynesia, there are a large number of flights between small regional airports on different islands but few outside of the community. Among the MSA's which represented the largest number of airports, were Jakarta (Indonesia), Auckland (New Zealand), Anchorage (Alaska, USA) and Port Moresby (Papua New Guinea) which are all regional hubs within sparsely-populated or archipelagic areas which may only feasibly be navigated by air. Since these small regional airports all agglomerate to a single MSA, it appears that a large volume of flights appear to leave from and arrive at the same MSA. Therefore these hubs have large unadjusted localities represented by large values along the diagonal. This artifact of the agglomeration process has a negligible affect on the structure of the network as a whole since these communities are not particularly central; this can be seen by the low centralities of these MSA's

Table 1: Table of regional hub MSA's and centralities

MSA	Centrality	Centrality/Centrality <sub>equal</sub>
Jakarta	0.00572	1.24
Auckland	0.00289	0.63
Anchorage	0.00092	0.20
Port Moresby	0.00176	0.38
Equal	0.0046	

The adjusted adjacency matrix used in the simulations and shown in Figure (2) maintains the strong continental community effect, however the localities have been uniformly set to 0.39 and a greediness of 30% has been applied. In a few cases this greediness leads to increased locality if all outgoing edges from an MSA move the message away from the nearest Tag city. The greediness also leads to a number of strong connections to Europe (but not directly the European Tag cities); the turquoise dots representing strength 0.3 in the columns on the right of the figure.

Table (2) shows the centralities of the most central MSAs in the network along with the Tag cities for comparison. All of the tag cities are above the baseline of equal centrality amongst all the nodes, however London, Washington DC and NY are especially so.

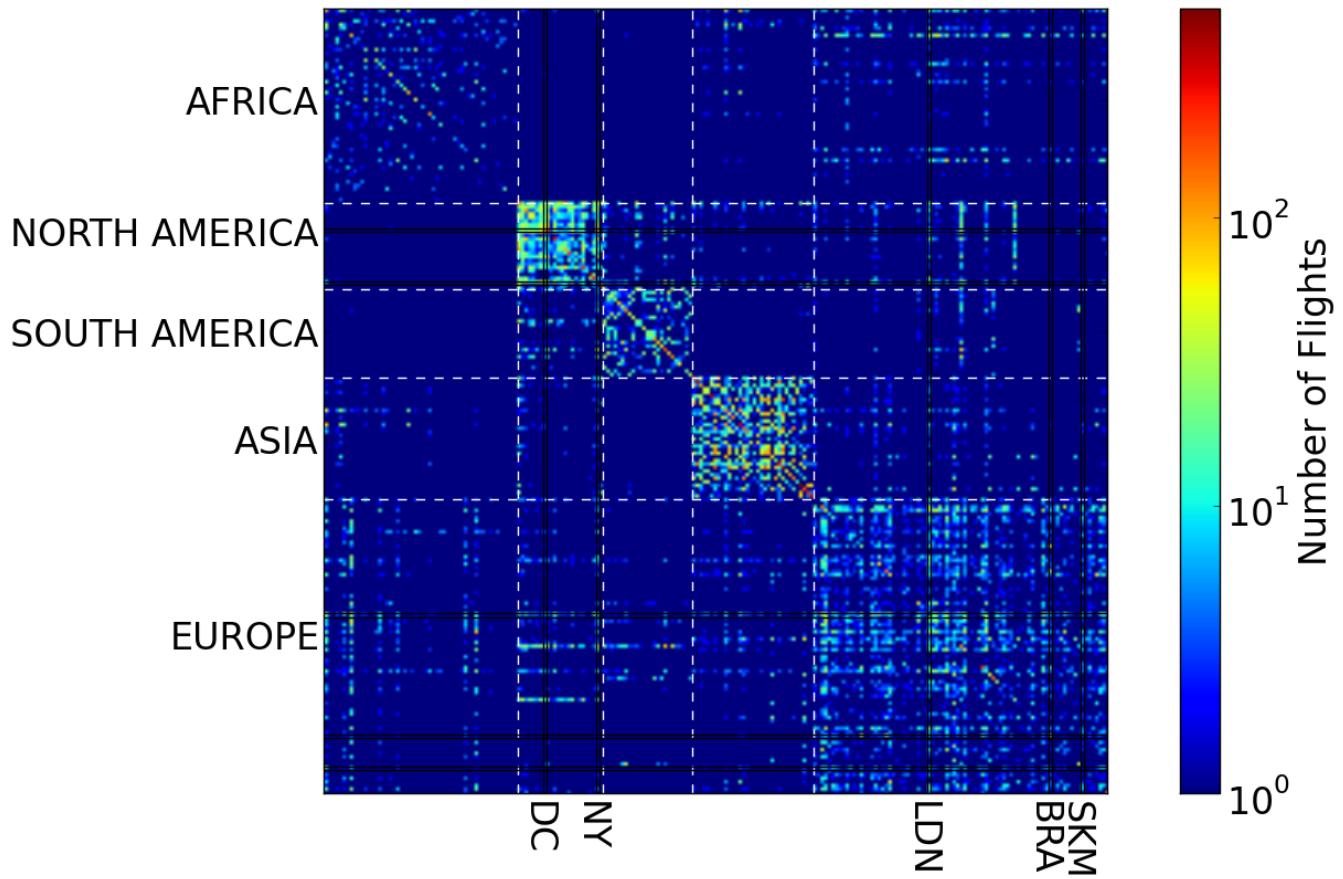


Figure 1: Heat map of raw flight numbers. Continent limits are marked by white dashed lines and tag cities with black lines.

Table 2: Table of MSA's with highest centrality values after locality adjustment (and tag cities for comparison)

MSA	Centrality	Centrality/Centrality <sub>equal</sub>
Shanghai	0.02243	4.87
London	0.02236	4.86
Chongqing	0.01953	4.25
Beijing	0.01688	3.67
LA	0.01677	3.65
Atlanta	0.01649	3.59
London	0.02236	4.86
NY	0.01597	3.47
DC	0.00976	2.12
Bratislava/Vienna	0.0066	1.43
Stockholm	0.0059	1.28
Equal	0.0046	

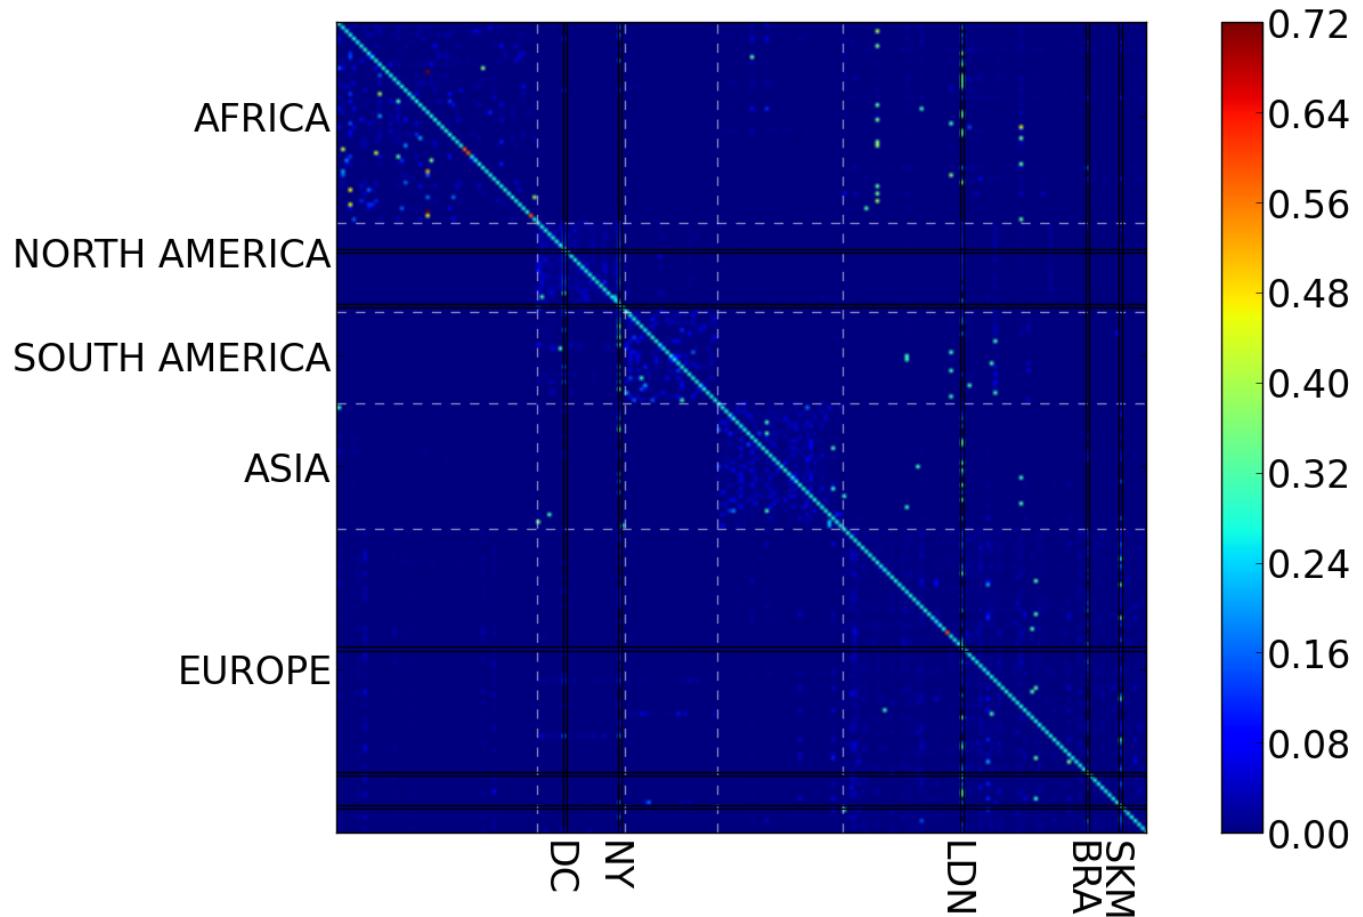


Figure 2: Heat map of normalised and locality-adjusted adjacency matrix with greediness set to 0.3. Continent limits are marked by white dashed lines and tag cities with black lines.

## B Reduced Network

Figure (3) shows the effective centralities of the cities within a reduced network comprising the cities of North America and Europe only (compare with the full global network shown in Fig (6) in the main paper). A degree of targeting of 30% now leads to a proportion of messages reaching the tag cities of 0.51 (compared to 0.46 using the full network). As expected the proportion of time spent in the tag cities increases as nodes are removed from the network. In fact the effect of the removal of the South American, African and Asian MSA's becomes smaller as targeting becomes stronger and routes the message towards the western hemisphere. Considering the pure, non-targeting random walk the reduced network increases the Tag proportion from 0.06 to 0.1; a percentage increase of 66%. However as the targeting becomes stronger this percentage difference becomes smaller. When greediness is set to 30% the reduced network increases the tag proportion from 0.45 to 0.51, an increase of only 13%.

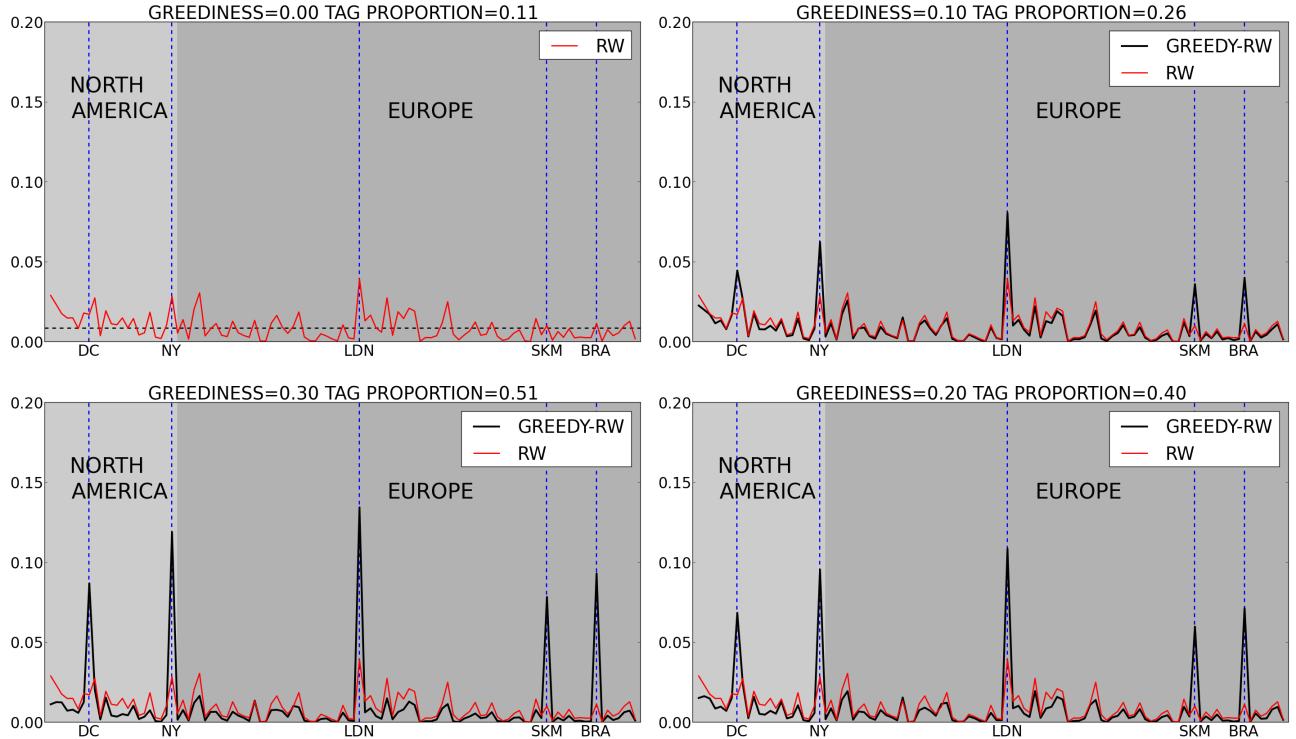


Figure 3: Plot of stationary distribution during a random walk on reduced MSA network (comprising only North America and Europe), with increasing degree of greediness moving clockwise from top left. The red line represents an unbiased random walk, corresponding to pure random mobilization via broadcast messaging. (Top left) The horizontal dashed line represents the uniform distribution of centralities expected in a fully connected graph. The black line in other plots represents a greedy random walk.

## C Website Traffic

Figure (4) shows the geographical distribution of traffic to our team's website in the 48 hours approaching the challenge. Traffic overwhelmingly originates from Europe and North America, particularly since this snapshot is from the critical latter stages of the propagation process, but we can also notice the presence of traffic originating from South America, Australia and Asia Pacific. The fact that tag traffic is significant even outside the Anglosphere suggests that the information diffusion either took place in languages other than English (a small but significant number of tweets were in languages other than English) or the English language media exposure was accessible via the *lingua franca*. While the South American, Asian and African nodes clearly participated in the diffusion, the network upon which this took place is likely somewhere between the reduced network presented here and the full global network presented in the main paper. Regardless of which extreme of network substrate dominates, we can conclude that significant targeting is required to reproduce the proportions of traffic towards the Tag cities.

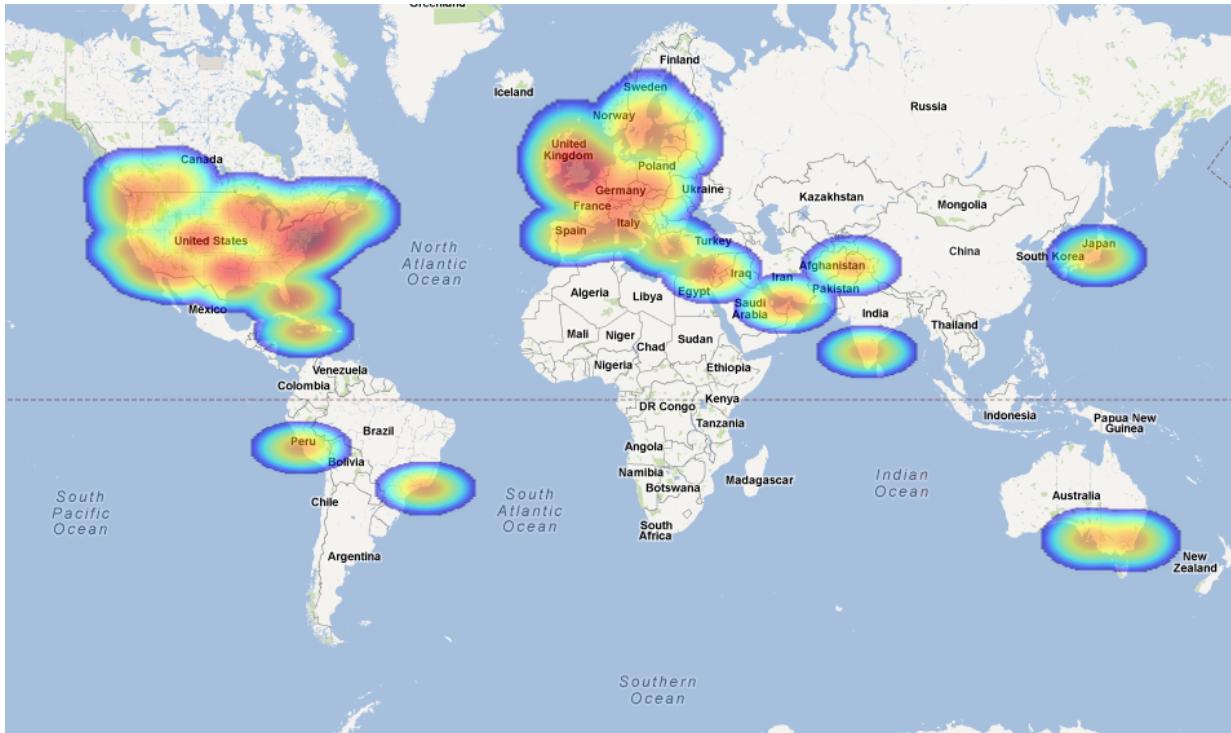


Figure 4: Heatmap showing traffic to crowdscanner.com on 48 hours approaching the challenge.